

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application of: Martin M. Deneroff et al.

MODULAR COMPUTING ARCHITECTURE HAVING COMMON COMMUNICATION
INTERFACE

Patent Docket No.: 499.034US1

09/29/99
JCS11 U.S. PTO
09/408874

PATENT APPLICATION TRANSMITTAL

BOX PATENT APPLICATION

Assistant Commissioner for Patents
Washington, D.C. 20231

We are transmitting herewith the following attached items and information (as indicated with an "X"):

- X Utility Patent Application under 37 CFR § 1.53(b) comprising:
X Specification (19 pgs, including claims numbered 1 through 31 and a 1 page Abstract).
X Formal Drawing(s) (18 sheets).
X Unsigned Combined Declaration and Power of Attorney (4 pgs).
X Return postcard.

The filing fee (NOT ENCLOSED) will be calculated as follows:

	No. Filed	No. Extra	Rate	Fee
TOTAL CLAIMS	31 - 20 =	11	x 18 =	\$198.00
INDEPENDENT CLAIMS	4 - 3 =	1	x 78 =	\$78.00
] MULTIPLE DEPENDENT CLAIMS PRESENTED				\$0.00
BASIC FEE				\$760.00
TOTAL				\$1,036.00

THE FILING FEE WILL BE PAID UPON RECEIPT OF THE NOTICE TO FILE MISSING PARTS.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.
P.O. Box 2938, Minneapolis, MN 55402 (612-373-6900)

By: Kent J. Sieffert
Atty: Kent J. Sieffert
Reg. No. 41,312

Customer Number **21186**

"Express Mail" mailing label number: EL334 875885US

Date of Deposit: September 29, 1999

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

By: Chris Hammond

Signature: Chris Hammond

MODULAR COMPUTING ARCHITECTURE HAVING COMMON COMMUNICATION INTERFACE

5

TECHNICAL FIELD

This invention is generally related to the field of high-speed computing systems and, more particularly, to a computer architecture having an flexible and incrementally scalable common communication interface that facilitates a wide variety of modular computing topologies.

10

BACKGROUND

With modern high performance computers, commonly known as supercomputers, there is an ever pressing need for more computing resources such as processors and input/output ports. Industry, therefore, is continuously developing high-speed computing systems in a wide variety of computing topologies and that support an increased number of computing resources. It is often difficult, however, to incrementally scale an existing computing system to a larger topology having more computing resources without disassembling the current configuration. In addition, it is often difficult to produce a wide variety of computing systems, including both low-end supercomputers and mid-range supercomputers, from a single manufacturing line. Thus, there is a need in the art for computer architecture in which high-speed computing systems can easily be configured and incrementally scaled in a modular fashion.

SUMMARY OF THE INVENTION

25

The present invention is directed to a distributed, shared memory computer architecture that is organized into a number of nodes, where each node has at least one processor. According to the invention, each node includes a common

communication interface that facilitates the ability to incrementally build and swap the nodes of the system without disrupting the overall computing resources of the system. More specifically, the common communication interface within each node connects local memory to the local processors, provides a port for communicating with a system-wide routing network and a port for communicating with an I/O subsystem. In this manner, each computing topology is a superset of smaller topologies supported by the architecture. As such, computing systems based upon the architecture may be easily and incrementally scaled without reconfiguring the existing components.

BRIEF DESCRIPTION OF THE DRAWING

Figures 1A and 1B are logical block diagrams that illustrate various embodiments of how the functionally independent modules of the inventive architecture may be combined in various topologies to form a high-speed computing systems.

Figure 1C illustrates three high-speed computing systems constructed using the flexible and scalable modular system described herein.

Figure 1D illustrates a high-performance processing system 140 having six vertical racks suitable for the modules described herein.

Figure 2A illustrates one embodiment of a C-Brick. This module contains four CPUs and eight memory slots.

Figure 3 illustrates an isometric view of a router module 300 herein referred to as the R-Brick.

Figure 4 illustrates an isometric view of an I/O module referred to herein as an X-Brick.

Figure 5 is an isometric view of an input/output module 500 herein referred to as an I-Brick.

Figure 6 illustrates one embodiment of an output module 600 referred to herein as a P-Brick, which is a more powerful input/output module than the I-Brick 500.

Figure 7 illustrates a computer rack 700 suitable for receiving one or more independent modules as described in this application.

Figure 8 is an isometric view illustrating a power module 800 herein referred to as a P-Bay.

5 Figure 9 illustrates a wiring diagram for a typical high-speed computing system comprising the flexible and scalable independently functional modules described herein.

Figures 10, 11 and 12 illustrate possible topologies based on the above discussed interconnect rules.

10 Figure 13 illustrates one embodiment of a common communication interface present within each node of the architecture, thereby facilitating their flexible and scalable interconnection.

Figure 14 illustrates one embodiment of an internal message format used by the common communication interface of Figure 13.

15 Figure 15 is a block diagram illustrating one embodiment of an internal, high-speed crossbar of the common communication interface.

DETAILED DESCRIPTION

The shortcomings, disadvantages and problems described in the background are addressed by the present invention, which will be understood by reading and
20 studying the following specification. The present invention is directed to a distributed, shared memory computer architecture that is organized into a number of nodes, where each node has at least one processor. According to the invention, each node includes a common communication interface that facilitates the ability to incrementally build and swap the nodes of the system without disrupting the overall
25 computing resources of the system. More specifically, the common communication interface within each node connects local memory to the local processors, provides a port for communicating with a system-wide routing network and a port for communicating with an I/O subsystem.

As described in detail below, independent routing modules are used within
30 the architecture to communicatively couple the nodes in a wide variety of

topologies. As a result, highly flexible and scalable high-speed computers may easily be constructed. The computer architecture is especially useful in constructing computing systems having a large number of processors, such as up to 4096, that share a single address space and require cache coherence across all of the nodes of the high-speed computer.

In this manner, these processing nodes and the routing modules are the basic building blocks for configuring a high-speed computing system and, therefore, are collectively referred to as bricks that may readily be interconnected in a variety of topologies. The computing system may include an arbitrary combination of processing nodes and other modules such that there need not be a fixed relation between the number of processing nodes and the other modules. Furthermore, as explained in detail below, each topology supported by the architecture is a superset of the smaller topologies supported by the architecture. Table 1 illustrates the various modules used within the computer architecture.

<u>Brick</u>	<u>Description</u>
C-Brick-MIPS	A CPU node populated with the MIPS processors.
C-Brick-Merced	A CPU brick populated Merced processors.
P-Brick	An IO brick that provides 14 PCI slots.
I-Brick	An I/O brick that supports complete IO needs of an entry level system.
X-Brick	An IO brick that provides four XIO slots.
R-Brick	A routing node providing eight routing ports.
Power Bay	A power bay providing power to a given rack of modules.
D-Brick	A diskette or disc drive module also referred to as a disc box.

Table 1

Figures 1A and 1B are logical block diagrams that illustrate various embodiments of how the functionally independent bricks described above may be combined in various topologies to form high-speed computing systems. For example, Figure 1A illustrates three computer topologies 10, 15 and 20. Topology 10 illustrates a single node having four processors (P) connected by a single

common communication interface (BR). Topology 20 illustrates an eight processor topology in which two nodes are communicatively coupled via their respective common communication interfaces. Topology 30 illustrates a sixteen processor topology in which four nodes are communicatively coupled via a single routing module. Figure 1B illustrates an extended hypercube computer topology in which nodes having 512 processors are interconnected according to the flexible architecture described herein.

As illustrated in Figures 1A and 1B, each computing topology is a superset of smaller topologies supported by the architecture. More specifically, topology 15 is a superset that includes two topologies 10. Similarly, topology 20 is a superset that includes four topologies 10, or two topologies 15. As such, computing systems based upon the inventive architecture described herein may be easily and incrementally scaled without reconfiguring the existing components.

Figure 1C illustrates an isometric view of three high-speed, rack-mounted computing systems constructed using functionally independent bricks and architecture of the present invention. More specifically, high-speed computing system 100 illustrates a low-cost entry into high-speed computing and includes power bay modules (Pwr Bay) 102, two processing nodes (C-Bricks) 104 that include multiple processors, a single routing module (R-Brick) 106 and a single input/output module (I-Brick) 108. High-speed computing system 120, on the other hand, illustrates an intermediate system with two processing nodes (C-Brick 124), a single routing module (R-Brick) 126, a single I/O module (I-Brick) 128, two power modules (Pwr Bay) 122 and three drive bays (D-Bricks) 129. High-speed computing system 130 illustrates a higher end computer having power modules (Pwr Bay) 132 and network connection module 133, four processing nodes (C-Bricks) 134 and routing module (R-Brick) 136. The four remaining I/O modules can be any combination of P-Bricks, I-Bricks and X-Bricks.

Figure 1D illustrates one quarter of a high-performance processing system 140. Six vertical racks of the twenty-four are illustrated. In the illustrated embodiment, computing system 140 can have up to eight C-Bricks per rack, up to 16 racks for a total of 512 CPUs. In addition, system 140 can have two or three

routing modules per rack and multiple power bays per rack. Computing system 140 further includes multiple I/O racks 142 that may comprise one to eight P, I or X Bricks.

5 **Processing Nodes (C-Bricks)**

Figure 2 illustrates one embodiment of a processing node 200 of the present invention. The processing node 200, referred to as a C-Brick, is a functionally independent module containing four local CPUs, local memory and the associated electronics required to operate as an independent distributed module. A C-Brick provides the following: 1) a high-speed serial channel to communicate between a system controller and the brick, 2) a high-speed serial channel to communicate with an internal level one (L1) system controller, 3) an external high-speed serial console port for a serial channel to communicate with an L1 system controller in an I/O-Brick.

15 The illustrated embodiment contains four CPUs 215 and eight memory slots 222. Memory slots 222 are designed to accept DIMM modules which support two rows of ten SDRAM chips. Front-mounted fans 228 are removable from the front, redundant and hot-swappable. External connectors 230 at the rear of the C-Brick 200 provide connections for power, routing network, I/O, and Universal Serial Bus (USB). As described below, the USB connector is used for connection to an optional level two (L2) system controller in small systems that do not have a routing module.

Processing node 200 includes a common communications interface 235 that facilitates the ability to incrementally build and swap the nodes of the system without disrupting the overall computing resources of the system. More specifically, the common communications interface 235 within node 200 connects local memory present in slots 222 to local processors 215 and, as discussed in detail below, provides a intelligent, high-speed interface to connectors 230.

30 **Routing Modules (R-Bricks)**

Figure 3 is a block diagram illustrating a routing module 300 herein referred to as an R-Brick. The R-Brick provides the following: 1) a high-speed serial channel to communicate between an internal L1 system controller and an internal router ASIC, 2) a USB slave port to communicate with a level two (L2) system controller and a level three (L3) system controller, 3) a USB hub with five master ports.

R-Brick 300 contains a single router ASIC 310, power circuitry 312, interface connectors 320 and 322, a level two (L2) system controller 315, USB hub 325 and USP upstream connector 327. Connectors 320 and 322 are D-Net connectors that, according to the invention, allow various modules to easily be interconnected and interchanged. Four of the D-Net connectors 320 of R-Brick 300 carry L2 system controller USB signaling. Another four D-Net connectors 322 carry all router-to-router communications. Thus, according to this embodiment, four C-Bricks may be coupled to R-Brick 300 via connectors 320 while four other routing modules may be connected to R-Brick 300 via connectors 322. In another embodiment, router module 300 has only six ports.

Input/Output Modules (I-Brick, P-Brick and X-Brick)

Three I/O modules are provided by the inventive computer architecture: an I-Brick, an X-Brick and a P-Brick. The I-Brick is intended to provide all of the IO needs for a basic system or to provide the boot requirements of a larger, more complex system. The P-brick provides twelve 64 bit PCI slots. The X-Brick is an IO expansion brick that provides four half-height XIO slots.

Figure 4 illustrates an isometric view of an I/O module referred to herein as an X-Brick. The X-Brick provides the following functionality: 1) a serial channel to communicate with an internal L1 system controller within a C-Brick, and 2) a means for reading and reporting the population of the I/O cards. More specifically, X-Brick 400 includes four I/O cards 404 that plug horizontally from the rear of box 402. A single host interface card also plugs horizontally from the rear. A mid-plane PCA 410 mounted vertically in the center of X-Brick 400 contains a single X-Bridge ASIC 412 for controlling I/O.

Figure 5 is an isometric view of an input/output module 500 herein referred to as an I-Brick. The I-Brick and the P-Brick provides the following: 1) a serial channel to communicate with the internal L1 controller within a C-Brick, 2) a means for reading and reporting the population of PCI cards, 3) a means for controlling the powering of PCI slots, 4) a means for controlling and monitoring the status of the power bay. In the illustrated embodiment, I-Brick 500 includes six PCI input/output boards 512 and a CDROM/DVD player 510. Power board 522 houses a single X-Bridge ASIC for controlling I/O. A pair of XTalk I/O ports are located on the rear of I/O module 500 and connect I-Brick 500 to one or two C-Bricks. In one embodiment, I-Brick 500 contains two removable fiber channel hard drives and a single removable media drive 510. Power board 522 mounts horizontally from the front of I-Brick 500 and accepts 48 volts DC as input and generates the required DC voltages for the system board, the PCI slots 512, disc drives 510.

Figure 6 illustrates one embodiment of an output module 600 referred to herein as a P-Brick, which is a more powerful input/output module than the I-Brick 500. More specifically, P-Brick 600 houses three X-Bridge ASICs and provides 12 PCI slots 602.

Racks and Power Bays

Figure 7 is an isometric view of a computer rack 700 suitable for receiving one or more independent modules as described in this application. For example, rack 700 can receive one or more C-Bricks, P-Bricks, I-Bricks or X-Bricks. In this manner, a scalable computer may easily be manufactured. Other embodiments of rack 700 are possible, such as a short rack or a dual-column rack.

Figure 8 is an isometric view illustrating a power module 800 herein referred to as a Pwr-Bay. In one embodiment, Pwr-Bay 800 holds up to six power supplies that are single-phase AC input and 950 watts at 48 VDC output each. As illustrated in Figure 8, Pwr-Bay 800 includes eight connectors 804 on the rear of the module. These connectors carry the 48 VDC power along with monitoring signals. P-Bay 804 further includes eight serial interfaces for monitoring each power supply. The

distributed architecture described herein requires, in one embodiment, a 48 volt supply.

System control and Interconnection Topologies

5 The distributed architecture of the present invention has a three-level hierarchy for system management. The heart of the architecture is an L1 system controller (not illustrated) that exists within each brick of the inventive architecture except the D-brick. This controller includes a microcontroller, a system monitoring chip, a scan interface chip, plus a collection of serial EPROMs, bus expanders and
10 communication interfaces which are specific to the device it controls. The L1 system controller is responsible for power control and sequencing, environmental control and monitoring, initiation of reset, and the storage of identification and configuration information for its host brick. The L1 system controller also provides console/diagnostic and scanned interfaces to the user.

15 Figure 9 illustrates a wiring diagram for a typical high-speed computing system comprising the functional modules described herein. An L2 system controller 904 provides rack-level system control, i.e., there is one L2 system controller 904 for each rack having C-Bricks. Moreover, the L2 system controller 904 acts as a central communications clearinghouse for the rack and controls all of
20 the bricks in that particular rack and associated I/O racks. In one embodiment, each L2 system controller 904 is equipped with a touch screen display and ethernet and modem ports, and can be used as a central point of control for the system. A third level of control is the L3 system controller 906 that provides a central point of control for the entire system and, in one embodiment, is a standalone workstation or
25 laptop.

 Computing system 900 includes a first rack 902 that contains two routing modules 910 (R-Bricks) and eight processing nodes 914 (C-Bricks). Because each routing modules 910 has four ports, four C-Bricks may be coupled to the router. Each C-Brick 914 is connected to one of the routing modules 910 via a single high-
30 speed USB cable. Similarly, each routing module 914 is connected to the L2 system controller 904 via USB cables. Computing system 900 further includes a meta-

router 912. The first CPU rack 902 also includes a single ethernet hub 914 that connects to additional CPU racks via local network 920. In addition, the L3 system controller may connect to other computers via network 925.

The L1 system controller within each brick provides direct low-level communications and control for all of the functions within a brick. In most systems, it acts as a slave to upstream L2 system controller 904. In one embodiment, however, which are limited to a smaller number of processing modules C-Bricks, one L1 system controller may act as a master controller for the entire system if no L2 system controller 904 is present.

The modules as described herein and their interconnects provide a wide variety of possible topologies having several different communication paths. For example, an L3 system controller may communicate directly to the L1 system controller of a C-Brick. Because a C-Brick has a standard upstream USB port, it is possible for the L3 to interface directly to the system without an L2. If there is an L2, then the upstream USB port of the C-Brick will be made inaccessible because the L1 of the C-Brick will already be using that USB channel to communicate with an R-Brick L1. It is also possible for an L3 system controller to communicate directly to an L2 system controller. As discussed above, this is typically via a network connection through a network hub. Otherwise, an L3 system controller can be connected directly to an L2 with a cross over twisted pair cable. If a routing module is included in the computing system, a L2 system controller is required.

The L2 system controller will act as a USB host for the particular rack. As described above, the L1 controller of the routing module contains a USB hub which drives the USB signaling to the local L1 and out on the four ports which are connected to the C-Bricks. It is therefore possible for an R-Brick to be coupled directly to a C-Brick. An R-Brick has a USB hub whose downstream ports will be routed in a shielded pair to the L1 system controller of the C-Brick. In some systems, it is possible for a C-Brick to communicate directly to another C-Brick. For example, if a routing module is not present and an L2 and L3 system controller is not present, then USB cannot be used as a communication mechanism. Thus, in a system with up to two C-Bricks, the L1 system controller of the bricks communicate

with each other via RS422 over cables. In addition, it is possible for the L1 system controller of a C-Brick to be coupled directly to the L1 system controller of an I/O-Brick. Since the C-Bricks and I/O-Bricks need to communicate in routerless systems, they are configured to communicate via RS422 over cables as well.

5 Figures 10, 11 and 12 illustrate possible topologies based on the above discussed possible interconnections. Interconnections 1005 illustrate computing systems that do not have a routing module and therefore C-Bricks are connected directly to I-Bricks. As illustrated, these interconnects use a serial 422 connection over a standard cable. Topologies 1010 illustrate computing systems having an L3
10 system controller and no L2 system controller. Here the L3 system controller uses a USB connection directly to the C-Bricks.

 In Figure 11, interconnections 1115 illustrate computing systems having three or more C-Bricks which, therefore, require a routing module, in which case an L2 system host is required. The L2 system controller is a USB host and the routing
15 modules are USB hubs. In Figure 12, interconnection topology 1208 illustrates a computing system having multiple L2 system controllers connected via a 10-base T hub. Here, an L3 system controller is optionally connected to the 10-base T hub.

Common Communication Interface

20 Figure 13 illustrates one embodiment of an inventive high-speed common communication interface 1300 for interconnecting the various independent modules described herein. According to the invention, common communication interface 1300 provides connectivity between the various modules in a fair and efficient manner. Each node within the computing system includes a common
25 communication interface 1300 that, as discussed in detail below, extends a "virtual" system bus throughout the distributed modules of the high-speed computing system. In this manner, processing nodes and other modules may be easily added and removed from the computing system.

 Common communication interface 1300 includes four distinct interfaces. A
30 A processor interface 1305 interfaces to one or more processing modules. Memory interface 1310 for interfacing to a portion of the global memory and for maintaining

cache coherency across the computing system. I/O interface 1315 for communicating with an I/O subsystem. Common communication interface 1300 further includes a router interface 1320 for interfacing to a router module.

Internally, common communication interface 1300 includes five interface
5 control units for managing its interfaces. More specifically, common communication interface 1300 includes a processor interface unit 1325, a memory interface unit 1330, an I/O interface unit 1335, a network interface unit 1340, and a local block 1345 for interfacing with the local chip resources. In this manner, common communication interface 1300 provides standard connectivity between
10 four types of external interfaces and an interface to local chip resources.

The interface control units of node 1300 are connected by a central crossbar 1350 for exchanging data between the interfaces at high data rates. In this manner, common communication interface 1300 facilitates distributed modular computing systems that share a single address space. In one embodiment, common
15 communication interface 1300 supports up to 256 processing nodes which, in one embodiment, comprise up to four processors each. Each interface control unit within common communication interface 1300 communicates by forwarding messages through, and receiving sending messages from, crossbar 1350. The messages used by the modules conform to a packetized network protocol. In one
20 embodiment, two types of messages are supported: requests and replies. This configuration helps the computing system avoid system deadlock situations and promotes cache coherence. When a message arrives through I/O interface 1335 or network interface 1330, the message is converted to an internal format. The reverse occurs when a message is sent via one of these interfaces.

25 The internal message format for common communication interface 1300 consists of a header frame that is a group of bits that is conceptually and logically a single unit. This header frame is optionally followed by one or more data frames carrying a total of 64-1,024 bits of data for the message. As each frame is received by or transmitted from common communication interface 1300. Control signals
30 embedded within the frame indicate all or some of the following information: 1) to

which interface control unit the frame is destined, 2) whether the frame is a request or reply, and 3) whether the frame concludes the current message.

Figure 14 illustrates one embodiment of an internal message format for common communication interface 1300. More specifically, Figure 14 illustrates control bits, header format and data formats. Within the header, a source is indicated via the most significant 11 bits. In one embodiment, the source identifies a device and a node.

Crossbar 1350 of common communication interface 1300 supports the flow of messages in the internal format discussed above along two virtual channels, multiplexed across physical channels connecting each unit of common communication interface 1300 to crossbar 1350. Crossbar 1350 is designed for minimal latency under light loads by means of a buffer queue bypass paths and maximum throughput under heavy loads by means of per virtual channel arbitration requests. It is in this manner that a plurality of nodes 135 may be used to interconnect the computing modules discussed herein in a variety of topologies.

Figure 15 is a block diagram of one embodiment of crossbar 1350 of common communication interface 1300. In the figure, a dual-FIFO refers to two virtual channel FIFOs within a single buffer memory structure. A quad-FIFO refers to four virtual channels in an analogous structure. Data path crossbar 1505 contains an 8-input by 6-output crossbar. The crossbar data path is 67 bits wide for all inputs and outputs and provides a 1.6 GB/s of data bandwidth per port at a 5 NS clock. The output queues provide buffering for outgoing unit messages and arbitrate for data path resources. The input queues provide buffering for data that has traversed the crossbar 1350 but has not yet been processed by its destination unit. Their primary role, therefore, is to provide rate matching and synchronization between crossbar 1350 and the receiving unit. Arbiter 1510 provides low-latency arbitration for uncontested ports via bypass arbitration and efficient high utilization via wave front arbitration as resources become saturated.

We claim:

1. A modular computing system comprising :
a set of functionally independent processing nodes having one or more local
5 processors and local memory, wherein each processing node includes a common
communication interface for communicating with other nodes within the system via
a messages conforming to a packetized network protocol; and
one or more routing modules communicatively coupling the processing
nodes via their respective common communication interface.
10
2. The system of claim 1, wherein the nodes operate in a global shared-memory
address space.
3. The system of claim 1, wherein the common communication interface within
15 each node connects the local memory to the local processors, provides at least one
port for interfacing with the routing modules and at least one port for
communicating with an input/output (I/O) subsystem.
4. The system of claim 1, wherein the common communication interface of
20 each node may be directly coupled together, thereby eliminating the routing module.
5. The system of claim 1, wherein the computing system includes a system
control hierarchy comprising a level one controller within each node, a level two
controller providing rack-wide control and a level three controller providing system-
25 wide control.
6. The system of claim 5, wherein the level one controller within each node
controls direct low-level communications within the node.
- 30 7. The system of claim 5, wherein each routing module comprises a level two
controller

8. The system of claim 5, wherein the level three controller is a standalone workstation.

5 9. The system of claim 1, wherein the nodes are communicatively coupled to the routing modules via a high-speed Universal Serial Bus.

10. The system of claim 1, wherein each common communication interface includes:

- 10 a processor interface for interfacing to one or more processing nodes.
a memory interface for interfacing to local memory as a portion of the global memory and for maintaining cache coherency across the computing system.
an I/O interface for communicating with an I/O subsystem.

15 11. The system of claim 1, wherein the common communication interface includes a plurality of interface control units.

12. The system of claim 11, wherein the common communication interface includes a central crossbar communicatively coupling each interface control unit for
20 the exchange of data between the external interfaces at high data rates.

13. The system of claim 12, wherein each interface control unit within common communication interface communicates by sending messages through the crossbar.

25 14. The system of claim 1, wherein the message protocol is a synchronous message protocol comprising requests and replies.

15. The system of claim 12, wherein the crossbar converts the messages to an internal message format.

30

16. The system of claim 12, wherein crossbar communicates the messages across two internal two virtual channels by multiplexing the messages across physical channels connecting each unit.

- 5 17. A processing node for a modular computing system comprising :
one or more local processors;
local memory;
a common communication interface coupled to the local processors and the
local memory, wherein the common communication interface includes:
10 a processor interface for communicating with one or more external
processing nodes;
a memory interface by which the local processors and the external
processor nodes communicating with the local memory;
a routing interface for communicating with an external routing
15 module; and
an I/O interface for communicating with an I/O external subsystem.

18. The processing node of claim 17, wherein the nodes operate in a global
shared-memory address space.

19. The processing node of claim 17, wherein the common communication
interface of the node may be directly coupled to a common communication interface
of another such node via the I/O interface.

20. The processing node of claim 5 and further including a system controller for
directing low-level communications within the node.

21. The processing node of claim 17, wherein the routing interface includes a
high-speed Universal Serial Bus.

22. The processing node of claim 17, wherein the common communication interface includes a plurality of interface control units.

23. The processing node of claim 22, wherein the common communication
5 interface includes a central crossbar communicatively coupling each interface control unit for the exchange of data between the external interfaces at high data rates.

24. The processing node of claim 22, wherein each interface control unit within
10 common communication interface communicates by sending messages through the crossbar.

25. The processing node of claim 17, wherein the message protocol is a
15 synchronous message protocol comprising requests and replies.

26. The processing node of claim 22, wherein the crossbar converts the messages to an internal message format.

27. The system of claim 22, wherein crossbar communicates the messages
20 across two internal two virtual channels by multiplexing the messages across physical channels connecting each unit.

28. A modular computing system comprising :
a set of functionally independent processing nodes operating in a global,
25 shared address space, wherein each node has one or more local processors and local memory, wherein each processing node includes a common communication interface for communicating with other modules within the system via a message protocol, and further wherein the common communication interface provides a single high-speed communications center within each node to operatively couple the
30 node to one or more external processing nodes, an external routing module, or an input/output (I/O) module.

29. The system of claim 28, wherein the computing system may include an arbitrary combination of processing nodes and other modules such that there need not be a fixed relation between the number of processing nodes and the other
5 modules.

30. A modular computing system comprising :
a set of functionally independent processing nodes that can be operatively coupled to form one of a plurality of computing topologies, wherein each computing
10 topology supports a number of the processing nodes, and further wherein each computing topology is a superset of the computing topologies that support fewer processing nodes.

31. The computing system of claim 28, wherein each processing node includes a
15 common communication interface for communicating with other processing nodes within the system

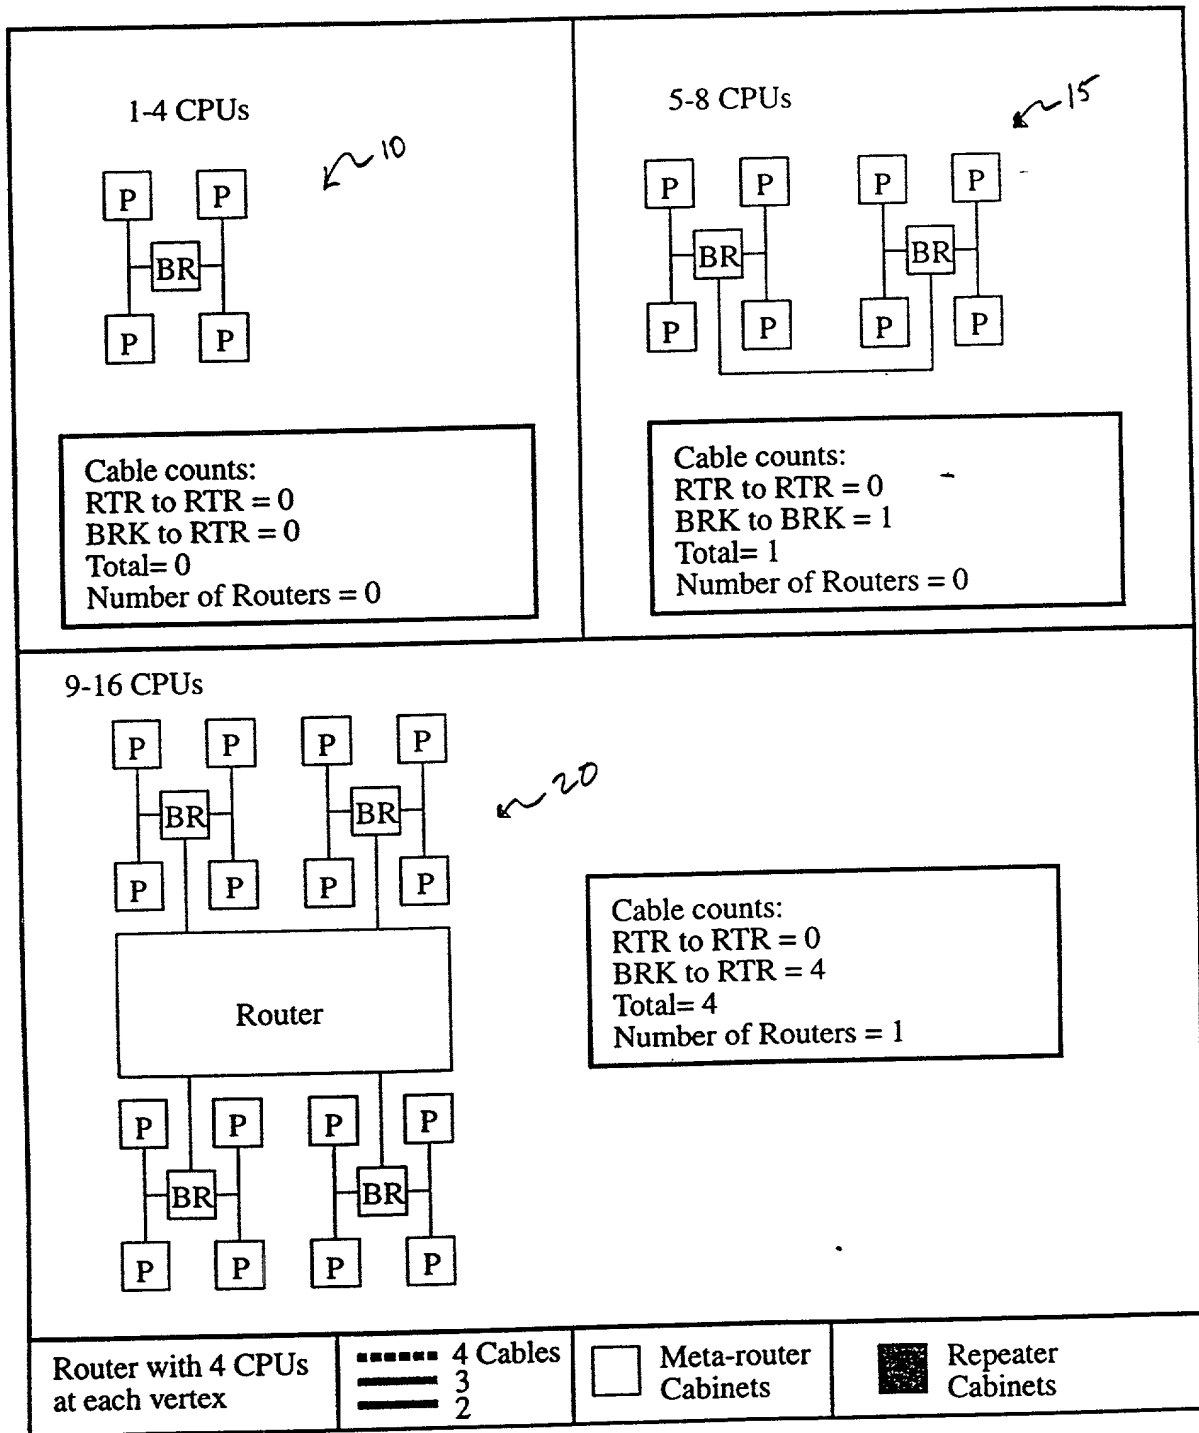
**MODULAR COMPUTING ARCHITECTURE HAVING COMMON
COMMUNICATION INTERFACE**


Abstract of the Disclosure

A distributed, shared memory computer architecture that is organized into a set of functionally independent processing nodes operating in a global, shared address space. Each node has one or more local processors, local memory and includes a common communication interface for communicating with other modules
5 within the system via a message protocol. The common communication interface provides a single high-speed communications center within each node to operatively couple the node to one or more external processing nodes, an external routing module, an input/output (I/O) module. The common communication interface that facilitates the ability to incrementally add and swap the nodes of the system without
10 disrupting the overall computing resources of the system.

"Express Mail" mailing label number: EL334 875885US
Date of Deposit: September 29, 1999
I hereby certify that this paper or fee is being deposited with the
United States Postal Service "Express Mail Post Office to Addressee"
service under 37 CFR 1.10 on the date indicated above and is
addressed to the Assistant Commissioner for Patents,
Washington, D.C. 20231
Printed Name Chris Hammond
Signature Chris Hammond

Fig 1A





512 CPU

Extended Hypercube Topology

FIG 1B








FIG 1C



Quadrant of 512P Single System Image

FIG 1D

FIG 2A



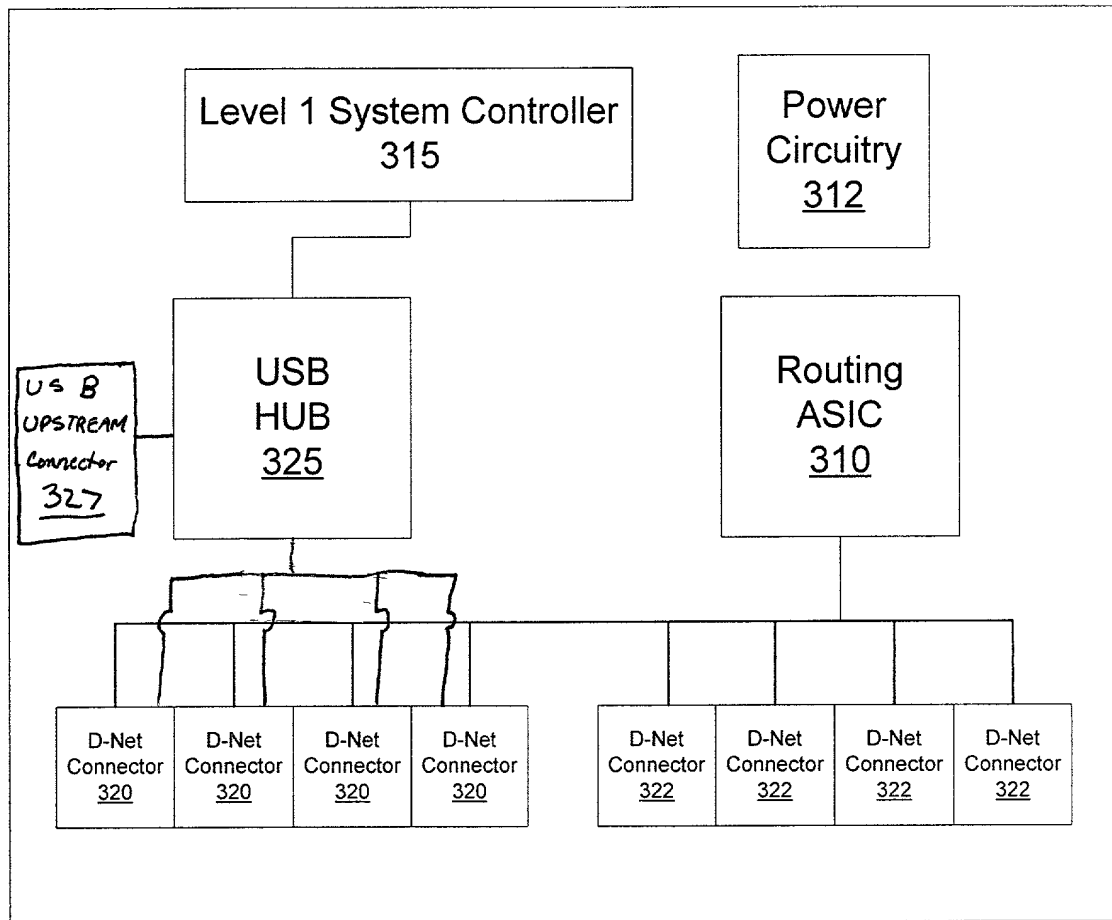


Figure 3






FIG. 21








FIG 7


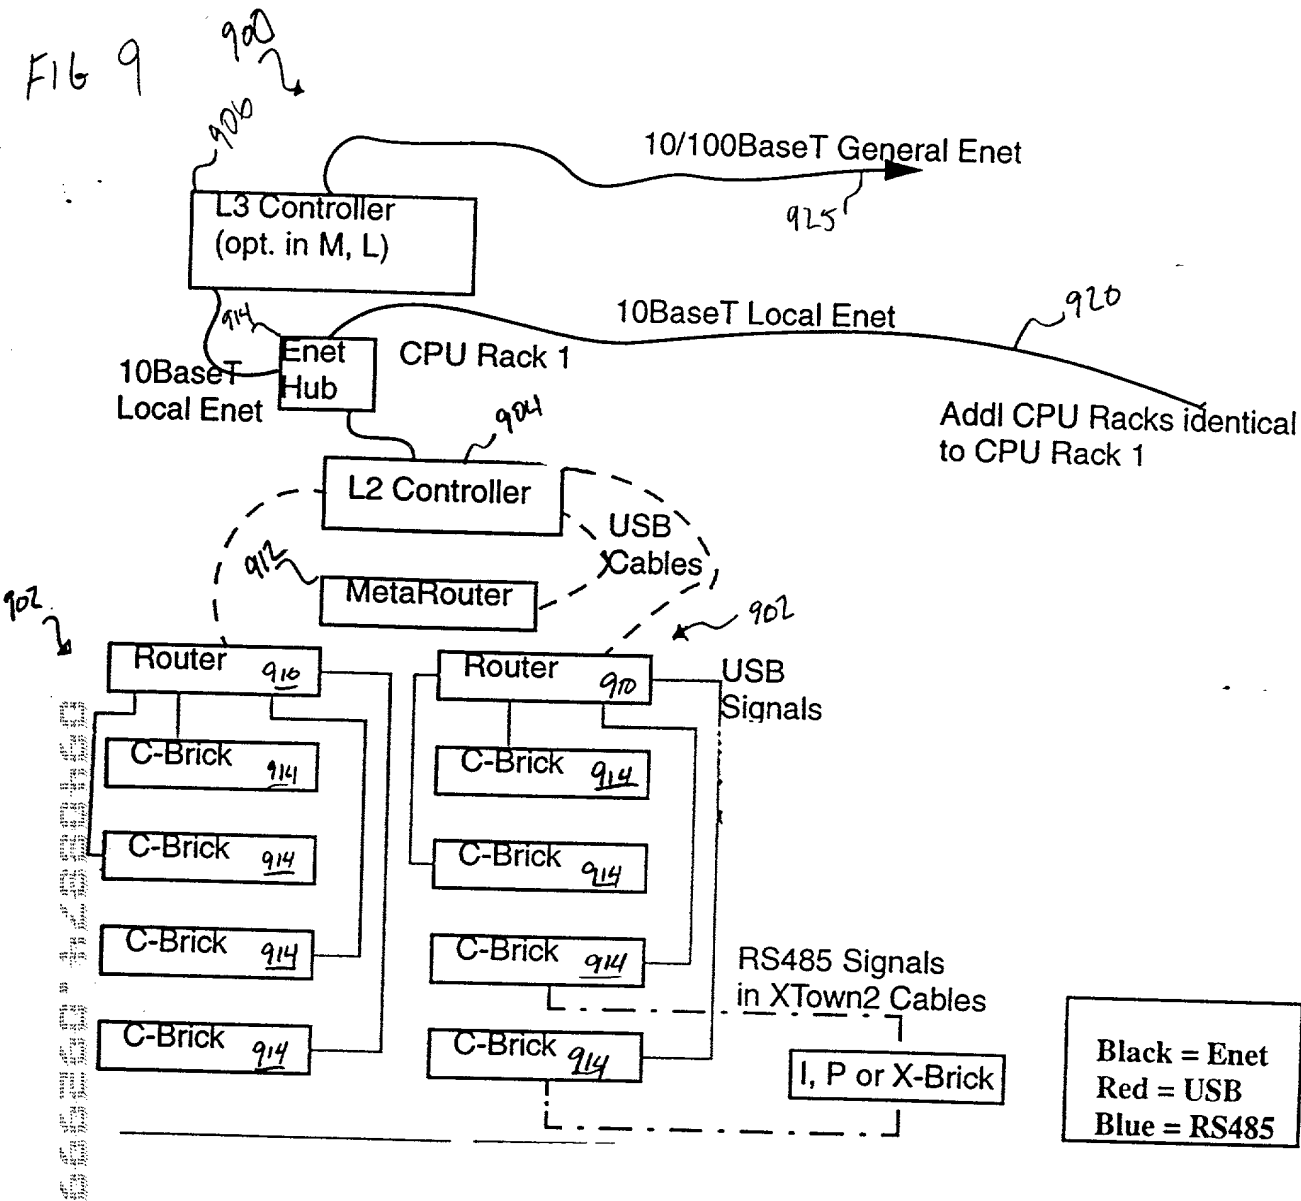





Fig 9



Year	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100
1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	



1010 ↗






FIG 10

1115 ~




FIG 11

FIG 12




FIG. 13






FIG 14


Control (5-bits)




Control (5-bits)




Header (67-bits)



Data (doubleword)




Data (quadword)



1350
7

FIG 15



SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

United States Patent Application

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor I hereby declare that: my residence, post office address and citizenship are as stated below next to my name; that

I verily believe I am the original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled: **MODULAR COMPUTING ARCHITECTURE HAVING COMMON COMMUNICATION INTERFACE.**

The specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with 37 C.F.R. § 1.56 (attached hereto). I also acknowledge my duty to disclose all information known to be material to patentability which became available between a filing date of a prior application and the national or PCT international filing date in the event this is a Continuation-In-Part application in accordance with 37 C.F.R. § 1.63(e).

I hereby claim foreign priority benefits under 35 U.S.C. § 119(a)-(d) or 365(b) of any foreign application(s) for patent or inventor's certificate, or 365(a) of any PCT international application which designated at least one country other than the United States of America, listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on the basis of which priority is claimed:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 119(e) of any United States provisional application(s) listed below:

No such claim for priority is being made at this time.

I hereby claim the benefit under 35 U.S.C. § 120 or 365(c) of any United States and PCT international application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT international application in the manner provided by the first paragraph of 35 U.S.C. § 112, I acknowledge the duty to disclose material information as defined in 37 C.F.R. § 1.56(a) which became available between the filing date of the prior application and the national or PCT international filing date of this application:

No such claim for priority is being made at this time.

I hereby appoint the following attorney(s) and/or patent agent(s) to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith:

Adams, Gregory J.	Reg. No. P-44,494	Forndenbacher, Paul J.	Reg. No. 42,546	Nama, Kash	Reg No 44,255
Adams, Matthew W.	Reg. No. 43,459	Forrest, Bradley A.	Reg. No. 30,837	Nelson, Albin J	Reg No 28,650
Anglin, J. Michael	Reg. No. 24,916	Harris, Robert J.	Reg. No. 37,346	Nielsen, Walter W	Reg No 25,539
Arora, Suneel	Reg. No. 42,267	Huebsch, Joseph C	Reg. No. 42,673	Oh, Allen J	Reg No 42,047
Bianchi, Timothy E.	Reg. No. 39,610	Jurkovich, Patti J.	Reg. No. P-44,813	Padys, Danny J	Reg No. 35,635
Billion, Richard E.	Reg. No. 32,836	Kalis, Janal M.	Reg. No. 37,650	Parker, J Kevin	Reg No 33,024
Black, David W.	Reg. No. 42,331	Kaufmann, John D.	Reg. No. 24,017	Peacock, Gregg A.	Reg No P-45,001
Brennan, Leoniede M.	Reg. No. 35,832	Klima-Silberg, Catherine I.	Reg. No. 40,052	Perdok, Monique M.	Reg. No 42,989
Brennan, Thomas F.	Reg. No. 35,075	Kluth, Daniel J.	Reg. No. 32,146	Polglaze, Daniel J	Reg No 39,801
Brooks, Edward J., III	Reg. No. 40,925	Lacy, Rodney L.	Reg. No. 41,136	Prout, William F.	Reg No 33,995
Byrne, Christopher	Reg. No. 32,204	Leffert, Thomas W.	Reg. No. 40,697	Schwegman, Micheal L	Reg. No. 25,816
Chu, Dinh C.P.	Reg. No. 41,676	Lemaire, Charles A.	Reg. No. 36,198	Sieffert, Kent J	Reg No 41,312
Clark, Barbara J.	Reg. No. 38,107	Litman, Mark A.	Reg. No. 26,390	Slifer, Russell D.	Reg No 39,838
Dahl, John M.	Reg. No. P-44,639	Lundberg, Steven W.	Reg. No. 30,568	Steffey, Charles E.	Reg No 25,179
Drake, Eduardo E.	Reg. No. 40,594	Mack, Lisa K.	Reg. No. 42,825	Terry, Kathleen R	Reg No 31,884
Eliseeva, Maria M.	Reg. No. 43,328	Maki, Peter C.	Reg. No. 42,832	Viksnins, Ann S.	Reg No 37,748
Embretson, Janet E.	Reg. No. 39,665	Malen, Peter L.	Reg. No. P-44,894	Weiner, Steve	Reg. No 38,330
Fernandez, Irene	Reg. No. 34,625	Mates, Robert E.	Reg. No. 35,271	Woessner, Warren D.	Reg No 30,440
Fogg, David N.	Reg. No. 35,138	McCrackin, Ann M.	Reg. No. 42,858		

I hereby authorize them to act and rely on instructions from and communicate directly with the person/assignee/attorney/firm/organization/who/which first sends/sent this case to them and by whom/which I hereby declare that I have consented after full disclosure to be represented unless/until I instruct Schwegman, Lundberg, Woessner & Kluth, P.A. to the contrary.

Please direct all correspondence in this case to **Schwegman, Lundberg, Woessner & Kluth, P.A.** at the address indicated below:

P.O. Box 2938, Minneapolis, MN 55402

Telephone No. (612)373-6900

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of joint inventor number 1 : **Martin M. Deneroff**

Citizenship: **United States of America**

Residence: **Palo Alto, CA**

Post Office Address: **2970 South Court Street
Palo Alto, CA 94306**

Signature: _____

Martin M. Deneroff

Date: _____

Full Name of joint inventor number 2 : **Steve Dean**

Citizenship: **United States of America**

Residence: **Mountain View, CA**

Post Office Address: **MS710
2011 North Shoreline Boulevard
Mountain View, CA 94043-1389**

Signature: _____

Steve Dean

Date: _____

X Additional inventors are being named on separately numbered sheets, attached hereto.

Serial No. not assigned

Filing Date: not assigned

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of joint inventor number 3 : **Timothy S. McCann**Citizenship: **United States of America**Residence: **Altoona, WI**Post Office Address: 419 Hampton Court
Altoona, WI 54720Signature: _____ Date: _____
Timothy S. McCannFull Name of joint inventor number 4 : **John Brennan**Citizenship: **United States of America**Residence: **Mountain View, CA**Post Office Address: MS710
2011 North Shoreline Blvd.
Mountain View, CA 94043-1389Signature: _____ Date: _____
John BrennanFull Name of joint inventor number 5 : **Dave Parry**Citizenship: **United States of America**Residence: **Mountain View, CA**Post Office Address: MS710
2011 North Shoreline Blvd.
Mountain View, CA 94043-1389Signature: _____ Date: _____
Dave ParryFull Name of joint inventor number 6 : **John Mashey**Citizenship: **United States of America**Residence: **Mountain View, CA**Post Office Address: MS710
2011 North Shoreline Blvd.
Mountain View, CA 94043-1389Signature: _____ Date: _____
John Mashey

§ 1.56 Duty to disclose information material to patentability.

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is canceled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is canceled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§ 1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

- (1) prior art cited in search reports of a foreign patent office in a counterpart application, and
- (2) the closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

- (1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or
- (2) It refutes, or is inconsistent with, a position the applicant takes in:
 - (i) Opposing an argument of unpatentability relied on by the Office, or
 - (ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

- (1) Each inventor named in the application;
- (2) Each attorney or agent who prepares or prosecutes the application; and
- (3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.